REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
1/20/2012	Final Report	8/1/2008 - 12/31/2011

4. TITLE AND SUBTITLE
Playing Games and Making Decisions with Complex State Spaces, Resource-Bound Agents, and Unforeseen Contingencies

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-08-1-0438

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Joseph Y. Halpern

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Cornell University, 341 Pine Tree Rd., Ithaca, NY 14850

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Sponsored Research, 875 N. Randolph St., Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2012-0286

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A - Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The goal of this project was to construct models for decision making and game theory that can deal with complex state spaces, resource-bounded and cognitively-limited agents, and unforeseen contingencies in a robust way. A number of advances were made in three areas:
1. A general game-theoretic framework for reasoning about strategic agents performing possibly costly computation was developed. The framework provides psychologically appealing explanations for behavior in well-studied games, such as (such as finitely repeated prisoner's dilemma and rock-paper-scissors).
2. A logic that can capture awareness and knowledge of lack of awareness was developed. The logic subsumes a number of logics introduced earlier. A sounds and complete axiomatization was provided for the logic.
3. A new solution concept, iterated regret mininmization, was defined and characterized. It was shown to predict behavior in a number of well-studied games.

15. SUBJECT TERMS
Game theory, decision theory, resource-bounded agents, awareness, regret miimization

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
U	U	U	SAR	10	Joseph Y. Halpern
					19b. TELEPHONE NUMBER *(Include area code)* 607-255-952

Reset

Playing Games and Making Decisions with Complex State Spaces, Resource-Bounded Agents, and Unforeseen Contingencies: Final Report

Grant No: AFOSR FA9550-08-1-0438
PI: Joseph Y. Halpern
Institution: Cornell University, Ithaca, NY 14853

Accomplishments

The goal of this project was to construct models for decision making and game theory that can deal with complex state spaces, resource-bounded and cognitively-limited agents, and unforeseen contingencies in a robust way.

A number of advances were made on this front, outlined below. Much of the work reported below is summarized in [4], which has been given as an invited talk at numerous conferences and universities.[1]

Taking computation into account in game theory: Nash equilibrium does not take computation into account. To see why this might be a problem, consider the following example, taken from [29,38].

Example 1: You are given an n-bit number x. You can guess whether it is prime, or play safe and say nothing. If you guess right, you get \$10; if you guess wrong, you lose \$10; if you play safe, you get \$1. There is only one Nash equilibrium in this 1-player game: giving the right answer. But if n is large, this is almost certainly not what people will do. Even though primality testing can be done in polynomial time, the costs for doing so (buying a larger computer, for example, or writing an appropriate program), will probably not be worth it for most people. The point here is that Nash equilibrium is not taking the cost of computing whether x is prime into account. ∎

There have been attempts in the game theory community to define solution concepts that take computation into account; we provide one that greatly generalizes an earlier approach due to Rubinstein. We consider *Bayesian games*, where each player has a *type*, which can be viewed as the player's private information. (In the example above, this would be the n-bit number.)

[1] A partial listing is given in the section on invited presentations below.

The players choose a Turing machine (TM) that plays for them. The input to the TM is the player's type. We associate with each TM and input a *complexity.* We then allow a player's utility to depend on the profile of moves played (the outputs of each player's TM), the profile of types, and the profile of complexities. Thus, we can "charge" for a TM that takes a long time to run, or uses a lot of space, on a given input. We can also charge for the complexity of the TM (i.e., the number of states in the TM); this corresponds to how complicated a player's program is. Moreover, we can take a given TM to be free, but charge for all others. This amounts to charging for searching for a different solution, and can be viewed as a way of capturing people's reluctance to switch from a strategy they are already comfortable with.

Formally, a *Bayesian machine game* G is a tuple $([m], \mathcal{M}, T, \Pr, \mathcal{C}_1, \ldots, \mathcal{C}_m, u_1, \ldots, u_m)$, where

- $[m] = \{1, \ldots, m\}$ is the set of players;
- $\mathcal{M}$ is the set of possible TMs;
- $\mathcal{T} \subseteq (\{0,1\}^*)^{m+1}$ is the set of type profiles, where the $(m+1)$st element in the profile corresponds to nature's type;
- Pr is a distribution on $\mathcal{T}$;
- $\mathcal{C}_i$ is a complexity function;
- $u_i : \mathcal{T} \times (\{0,1\}^*)^m \times I\!N^m \to I\!R$ is player i's utility function. Intuitively, $u_i(\vec{t}, \vec{a}, \vec{c})$ is the utility of player i if $\vec{t}$ is the type profile, $\vec{a}$ is the action profile (where we identify i's action with M_i's output), and $\vec{c}$ is the profile of machine complexities.

A *strategy* for player i is a function from T_i to $\Delta(A_i)$ (where, as usual, we denote by $\Delta(X)$ the set of distributions on the set X). If σ is a strategy for player i, $t \in T_i$ and $a \in A_i$, then $\sigma(t)(a)$ denotes the probability of action a according to the distribution on acts induced by $\sigma(t)$. Given a joint strategy $\vec{\sigma}$, we can take $u_i^{\vec{\sigma}}$ to be the random variable on the type space T defined by taking $u_i^{\vec{\sigma}}(\vec{t}) = \sum_{\vec{a} \in A} (\sigma_1(t_1)(a_1) \times \ldots \times \sigma_m(t_m)(a_m)) u_i(\vec{t}, \vec{a})$. Player i's expected utility if $\vec{\sigma}$ is played, denoted $U_i(\vec{\sigma})$, is then just $\mathbf{E}_{\Pr}[u_i^{\vec{\sigma}}] = \sum_{\vec{t} \in T} \Pr(\vec{t}) u_i^{\vec{\sigma}}(\vec{t})$.

With these definitions in hand, we can define *computational Nash equilibrium* in the standard way: a machine profile $\vec{M} \in \mathcal{M}^m$ is a Nash equilibrium if, for all players i and TMs $M_i' \in \mathcal{M}$,

$$U_i^G[(M_i, \vec{M}_{-i})] \geq U_i^G[(M_i', \vec{M}_{-i})].$$

Thus, for each player i, TM M_i is a best response to M_{-i}. We can show that, in general, a computational Nash equilibrium does not exist. This is because we do not allow mixing over TMs. While TMs can randomize, we want to allow "charging" for such. If we do not charge for randomization, then it can be shown that a computational Nash equilibrium exists, provided that the utility functions and the type distribution are computable [29,38].

In [29,38], the framework is used to give simple explanations of phenomena such as cooperation in repeated Prisoner's Dilemma. It is also shown that there are deep connections between computational games and cryptographic protocol security. In [41], the framework is applied to decision theory. It is shown that the approach can be used to explain well-known phenomena such as the *status quo bias* (people are much more likely to stick with what they already have) and *belief polarization* (after hearing the exactly same information, two people with possibly different prior beliefs can end up with diametrically opposed conclusions). The framework is also used to define two extensions of the standard notion of *value of information*: *value of computational information* and *value of conversation.*

Awareness: In game theory, it is typically assumed that the game (i.e., the agents who are playing, the moves that each one can make, and the payoffs at the leaves) is *common knowledge*: all the agents know them, they know they know, and so on. But this assumption of common knowledge is not reasonable in general. In large auctions on the internet, players may certainly not know how many other players there are in the game, let alone who they are. There has been a great deal of work on awareness in the AI and game theory literatures, going back to the original paper written by Fagin and Halpern (Belief, awareness, and limited reasoning, *Artificial Intelligence* **34**, 1988, pp. 39–76). None of the earlier work dealt with knowledge of lack of awareness (indeed, some of the approaches used in game theory are provably incapable of dealing with it). This is particularly important if we interpret lack of awareness as stemming from lack of computational power. An agent might well know that another might have more computational abilities than he does, and so know things that he is not aware of (because he cannot compute them). Rêgo and I described and provided a complete exiomation of a logic that could capture knowledge of lack of awareness [27], using quantification over propositions. The idea was that you could say "there is something I am not aware of, although I don't know what it is".

Unfortunately, the logic has a significant problem: it is impossible for an agent to be uncertain about whether he is aware of all formulas. This problem is corrected in [28]. The solution involves considering models where there are different languages associated with each world; thus, the agent has uncertainty about what the laguage is. This approach allows us to unify and extend much of the earlier work on awareness.

Regret minimization: Perhaps the most common solution concept considered in game theory is Nash equilibrium. While many other solution concepts have been considered in the literature, there are a number of well-known games where none of them seems appropriate.

Example 2: Consider the well-known *Traveler's Dilemma* (K. Basu, The traveler's dilemma: paradoxes of rationality in game theory, *American Economic Review* **84**:2, 1994, pp. 391–395): Suppose that two travelers have identical luggage, for which they both paid the same price. Their luggage is damaged (in an identical way) by an airline. The airline offers to recompense them for their luggage. They may ask for any dollar amount between \$2 and \$100. There is only one catch. If they ask for the same amount, then that is what they will both receive. However, if they ask for different amounts—say one asks for $\$m$ and the other for $\$m'$, with $m < m'$—then whoever asks for $\$m$ (the lower amount) will get $\$(m + 2)$, while the other traveler will get $\$(m - 2)$.

It seems at first blush that both travelers should ask for \$100, the maximum amount, for then they will both get that. However, one of them might then realize that he is actually better off asking for \$99 if the other traveler asks for \$100, since he then gets \$101. In fact, \$99 *weakly* *dominates* \$100, in that no matter what Traveler 1 asks for, Traveler 2 is always at least as well off asking for \$99 than \$100, and in one case (if Traveler 2 asks for \$100) Traveler 1 is strictly better off asking for \$99. Thus, it seems we can eliminate 100 as an amount to ask for. However, if we eliminate 100, a similar argument shows that 98 weakly dominates 99. And once we eliminate 99, then 97 weakly dominates 98. Continuing this argument, both travelers end up asking for \$2! In fact, it is easy to see that (2,2) is the only Nash equilibrium. With any other pair of requests, at least one of the travelers would want to change his request if he knew what the other traveler was asking. Nevertheless, in experiments, it has been shown that people typically ask for amounts in the high 90's; very few people ask for 2. ∎

Pass and I [28] introduce a new solution concept, *iterated regret minimization*, which has the same qualitative behavior as that observed in the experiments, not just in Traveler's Dilemma, but in many other games that have proved problematic for Nash equilibrium, including the Centipede Game, Nash bargaining, and Bertrand competition. As the name suggests, iterated regret minimization involves the iterated deletion of strategies that do not minimize *regret*. It is a solution concept that leads to quite different predictions than Nash equilibrium, while being grounded on regret, a decision rule that has a long history in decision theory, and has a great deal of support from the psychology literature.

Roughly speaking, the idea of regret in decision theory is that an agent chooses an action a that minimizes regret across states, where the regret of action a in a state s is the difference between the agent's utility when he performs a in a state s and when he performs the act that gives the highest utility in state s. To apply regret in a strategic-form game, we take the states to be the other players' strategy choices. Iterated regret minimization takes this idea one step further: we see what inferences we can draw starting with minimal beliefs about the other players' strategies, using only the fact that the other players are regret minimizer.

Personnel Supported

- Joseph Y. Halpern (PI)
- Adam Bjorndahl (Ph.D. student)
- Samantha (Yinyee) Leung (Ph.D. student)
- Vasumathi Raman (Ph.D. student)
- Nan Rong (Ph.D. student)
- Dongcai Shen (Ph.D. student)

New discoveries, inventions or patent disclosures

- D. J. Martin, J. Y. Halpern, and J. Gehrke, System and Method for Scalable Sponsored Auctions, patent application filed August, 2008.

Selected Invited Presentations

- Beyond Nash equilibrium: solution concepts for the 21st century
  - Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR 2008), Sydney, Australia (September, 2008) (invited talk).
  - GAMES Summer School, Bertinoro, Italy (June, 2009) (invited talk)
  - Israeli Game Theory Society Seminar (invited talk), Nov. 2009
  - Behavioral and Quantitative Game Theory: Conference on Future Directions, Newport Beach, California, May, 2010 (invited plenary talk)
  - Saul Gorn Memorial Lecture, University of Pennsylvania, April, 2011
  - ACM Conference on Autonomous Agents and Multiagent System (invited plenary talk), Taiwan, May, 2011
  - Conference on Computing in the 21st Century (keynote talk), Beijing, October, 2011
- Constructive Decision Theory
  - London School of Economics, January, 2011
  - European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECQARU), July, 2011 (invited presentation)
- Game Theory with Costly Computation
  - Center for Rationality, Hebrew University, October, 2009
  - Tel Aviv University Economics Dept., October, 2009 (Sackler Lecture)
  - Champ'ery Spring School, Champery, Switzerland, Feb., 2010
- Iterated Regret Minimization: A New Solution Concept
  - Statistics Dept., Tel Aviv University, October, 2009 (Sackler Lecture)
  - Center for Rationality, Hebrew University, Jan. 2010
  - Champ'ery Spring School, Champery, Switzerland, Feb., 2010
- Robustness and Optimization of Scrip Systems
  - Tel Aviv University, CS Dept., October, 2009 (Sackler Lecture)
  - Hebrew University, CS Dept., October, 2010
  - EPFL, Lausanne, February, 2010
  - Weizmann Institute, March, 2010
- Causality, responsibility, and blame: a structural-model approach
  - Workshop on Theory and Practice of Provenance, San Francisco (February, 2009) (invited talk)

  - MICRAC Workshop on Causality in AI and Cognitive Psychoology Toulouse, France (June 2009) (invited talk)
  - International Conference on Scalable Uncertainty Management, Dayton, Ohio (October 2011) (invited talk)

- Intransitivity and Vagueness
  - 31st Linz Seminar on Fuzzy Set Theory, Linz, Austra (February 2010) (invited talk)

- Knowledge and common knowledge in a distributed environment
  - Fulbright Distinguished Chair Lecture, Jerusalem (April 2010)

Honors/Awards

Awards during reporting period:

- selected IEEE Fellow, 2012
- Selected Economic Theory Fellow, Society for the Advancement of Economic Theory, 2011
- ACM SIGART Autonomous Agents Research Award, 2011
- Saul Gorn Memorial Lecturer, University of Pennsylvania, 2011.
- Fulbright Distinguished Chair in Natural Sciences and Engineering, Hebrew University, 2009-10
- Kenneth A. Goldman '71 Excellence in Teaching Award, 2010
- Sackler Lecturer, Tel Aviv University, 2009
- 2009 Edsger Dijkstra Prize in Distributed Computing
- 2008 ACM/AAAI Newell Award

Earlier significant awards:

- ACSU (Association of Computer Science Undergraduates) Faculty of the Year Award, 2007.
- "Redoing the Foundations of Decision Theory" given the Ray Reiter Best Paper award at the Tenth International Conference on Knowledge Representation and Reasoning, 2006.
- Selected Fellow of AAAS, November, 2005.
- Selected Fellow of ACM, 2002.
- Fulbright Fellow, 2001-02.
- Guggenheim Fellow, 2001-02.

- Milner Lecturer, University of Edinburgh, 2000.
- Awarded 1997 Gödel Prize for outstanding paper in the area of theoretical computer science for "Knowledge and common knowledge in a distributed environment".
- Fellow of the American Association of Artificial Intelligence, 1993.
- "An analysis of first-order logics of probability" given the Publisher's Prize as best paper of the 11th International Joint Conference on Artificial Intelligence, 1989.
- "Belief, awareness, and limited reasoning" given the MIT Press Publisher's Prize as best paper of the 9th International Joint Conference on Artificial Intelligence, 1985.

Publications

A. Books

1. R. Dechter, H. Geffner, and J.Y. Halpern (eds.), Heuristics, Probability and Causality: A Tribute to Judea Pearl, College Publications, 2010.

B. Articles in Books

2. J. Y. Halpern, Computer science and game theory: A brief survey, in The New Palgrave Dictionary of Economics (S. N. Durlauf and L.E. Blume, eds.), Palgrave MacMillan, 2008.

3. J. Y. Halpern and C. Hitchcock, Actual causation and the art of modeling, in "Heuristics, Probability and Causality: A Tribute to Judea Pearl" (editors, R. Dechter, H. Geffner, and J. Y. Halpern), College Publications, 2010, pp. 383-406.

4. J.Y. Halpern Beyond Nash equilibrium: solution concepts for the 21st century, in "Lectures in Game Theory for Computer Scientists" (K. R. Apt and E. Grädel, editors), Cambridge University Press, 2011, pp. 264–289.

5. J.Y. Halpern and R. Pass, That's all I know: A logical characterization of iterated admissibility, in "Knowing, Reasoning, and Acting: Essays in Honour of Hector J. Levesque" (G. Lakemeyer and S. A. McIlraith, editors), College Publications, 2011.

C. Journal Publications

6. J. Y. Halpern and V. Weissman, A formal foundation for XrML, Journal of the ACM 55:1, 2008.

7. J. Y. Halpern and V. Weissman, Using first-order logic to reason about policies, ACM Transactions on Information and System Security 11:4, 2008.

8. J.Y. Halpern and K. O'Neill, Secrecy in multi-agent systems, ACM Transactions on Information and System Security 12:1, 2008.

9. J.Y. Halpern, A nonstandard characterization of sequential equilibrium, perfect equilibrium, and proper equilibrium, International Journal of Game Theory 38:1, 2009, pp. 37-50.

10. J.Y. Halpern, Intransitivity and vagueness, Review of Symbolic Logic 1:4, 2009, pp. 530-547.

11. J.Y. Halpern and L.C. Rêgo, Reasoning about knowledge of unawareness, Games and Economic Behavior, 67:2, 2009, pp. 503-525.

12. J. Y. Halpern, D. Samet, and E. Segev, On definability in modal logic, Review of Symbolic Logic, 2:3, 2009, pp. 451-468.

13. J. Y. Halpern, D. Samet, and E. Segev, Defining knowledge in terms of belief: the modal logic perspective, Review of Symbolic Logic, 2:3, 2009, pp. 469-487.

14. J.Y. Halpern and R. Pucella, Evidence with uncertain likelihoods, Synthese, 171:1, 2009, pp. 111–133.

15. J.Y. Halpern, Lexicographic probability, conditional probability, and nonstandard probability, Games and Economic Behavior 68:1, 2010, pp. 155-179.

16. M. M. Halldorsson, J.Y. Halpern, L. Li, and V. Mirrokni, On spectrum sharing games, Distributed Computing 22:4, 2010, pp. 235-248.

17. J. Y. Halpern and S. Petride, A knowledge-based analysis of global function computation, Distributed Computing 23:3, 2010, pp. 197–224.

18. M. Goldstein, H. Westerfall, J. Y. Halpern, A. Lotem, L. Onnis, J. Schwade, S. Edelman, General cognitive principles for learning structure in time and space, Trends in Cognitive Science 14:6, 2010, pp. 249-258.

19. A. Meliou, W. Gatterbauer, J. Y. Halpern, C. Koch, K. F. Moore, and D. Suciu, Causality in Databases, IEEE Data Engineering Bulletin 33:3, 2010, pp. 59-67.

20. J. Y. Halpern and R. Pucella, Dealing with logical omniscience, Artificial Intelligence 175:1, 2011, pp. 220-235.

21. M. Bickford, R. L. Constable, J. Y. Halpern, and S. Petride, Knowledge-based synthesis of distributed systems using event structures, Logical Methods in Computer Science 7:2, 2011.

22. I. A. Kash, E. J. Friedman, and J.Y. Halpern, Multiagent learning in large anonymous games, Journal of AI Research 40, 2011, pp. 571-598.

23. D. Parkes and J.Y. Halpern, Viewpoint: Journals for Certification, Conferences for Rapid Dissemination, Communications of the ACM, 2011, pp. 36-38.

24. L.C. Rêgo and J.Y. Halpern, Generalized solution concepts in games with possibly unaware players, International Journal of Game Theory, 2011, pp. 1-25.

25. I.A. Kash, E.J. Friedman, and J.Y. Halpern, Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion, Distributed Computing, 2011, to appear.

26. P.D. Grunwald and J.Y. Halpern, Making decisions using sets of probabilities: updating, time consistency, and calibration, Journal of AI Reasoning 42, 2011, pp. 393-426.

27. J. Y. Halpern and R. Pucella, Modeling adversaries in a logic for security protocol analysis, Logical Methods in Computer Science, 2011.

28. J. Y. Halpern and R. Pass, Iterated regret minimization: a new solution concept, Games and Economic Behavior, 74:1, 2012, pp. 194–207.

29. J. Y. Halpern and R. Pass, Algorithmic rationality: game theory with costly computation, submitted.

D. Articles in Conference Proceedings

30. P.D. Grunwald and J.Y. Halpern, A game-theoretic analysis of updating sets of probabilities, Proceedings of the Twenty-Fourth Conference on Uncertainty in AI, 2008, pp. 240-247 (this is a preliminary version of #26).

31. D.J. Martin and J.Y. Halpern, Shared winner determination in sponsored search auctions, Proceedings of the 25th International Conference on Data Engineering, 2009, pp. 270-280.

32. I.A. Kash, E.J. Friedman, and J.Y. Halpern, Multiagent learning in large anonymous games, Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2009, pp. 765–772 (this is a preliminary version of #22).

33. I. A. Kash, E.J. Friedman, and J.Y. Halpern, Manipulating scrip systems: sybils and collusion, Proceedings of the First Conference on Auctions, Market Mechanisms, and Multiagent Systems (AMMA), 2009.

34. J.Y. Halpern and R. Pass, Iterated regret minimization: A more realistic solution concept, Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009), 2009, pp. 153–158 (this is a preliminary version of #28)

35. J.Y. Halpern and R. Pass, A logical characterization of iterated admissibility, Proceedings of Twelfth Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 2009, pp. 146-155 (this is a preliminary version of #5).

36. J.Y. Halpern, R. Pass, and V. Raman, An epistemic characterization of zero knowledge, Proceedings of Twelfth Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 2009, pp. 156-165.

37. J.Y. Halpern and L.C. Rêgo, Reasoning about knowledge of unawareness revisited, Proceedings of Twelfth Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 2009, pp. 166-173.

38. J.Y. Halpern and R. Pass, Game theory with costly computation: formulation and application to protocol security, First Symposium on Innovations in Computer Science, 2010, pp. 120-142.

39. J.Y. Halpern and N. Rong, Cooperative equilibrium, Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), 2010, pp. 1465-1466.

40. J.Y. Halpern, From causal models to counterfactual structures, Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR 2010), 2010.

41. J.Y. Halpern and R. Pass, I don't want to think about it now: Decision theory with costly computation, Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR 2010), 2010.

42. J.Y. Halpern and N. Rong, MDPs with Unawareness, Proceedings of the Twenty-Sixth Conference on Uncertainty in AI, 2010.

43. A. Bjorndahl, J.Y. Halpern, and R. Pass, Reasoning about justified belief, Proceedings of Thirteenth Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 2011, pp. 221–227.

44. D. Dolev, D. Feitelson, J.Y. Halpern, R. Kupferman, and N. Linial, No justified complaints: on fair sharing of multiple resources, Proceedings of 3rd Conference on Innovations in Theoretical Computer Science (ITCS 2012), 2012.

45. J.Y. Halpern and W. Kets, Ambiguous language and differences in beliefs, submitted.

46. J.Y. Halpern and S. Leung, Weighted sets of probabilities and minimax weighted expected regret: new approaches for representing uncertainty and making decisions, submitted.